

# Remote File Transfer Commands

## Category: File Transfers

### DRAFT

This article is being reviewed for completeness and technical accuracy.

The following file transfer commands can be used when the source and destination are located at different hosts. They can be used to transfer files either between NAS HECC hosts or between a NAS host and a remote host such as your local desktop system.

- **scp (with/without HPN-SSH patch)**

Secure Copy Protocol (SCP), based on Secure Shell Protocol (SSH), is a means of securely transferring files between a local and a remote host. Both the authentication information (such as password or passcode) and user's data are encrypted.

#### **Normal scp (without the HPN-SSH patch)**

The most widely used scp is from OpenSSH.

- ◆ *Where is it installed at NAS?*

A copy of scp from OpenSSH without the patch is available on the Pleiades front-end and bridge nodes (pfe[1-12], bridge[1-2]), all Columbia nodes, Lou[1-2], SFE[1,2], Bouncer, and Bruiser.

The copy on SUP contains the HPN-SSH patch.

scp is not available on DMZFS[1,2]. Use scp on Columbia, Pleiades, Lou or your localhost to push files into DMZFS[1,2] or pull files out of DMZFS[1,2].

- ◆ *Do you need it installed on your localhost?*

If you already have a version of SSH installed on your localhost, most likely, scp is already there.

- ◆ *When to use it?*

scp is typically used for transferring small files (<< 5GB) within NAS or offsite (<< 1 GB) that takes reasonable amount of time to complete.

◆ *Examples:*

For outbound transfer:

```
lou1% scp local_username@your_localhost.domain:foo ./foo2
```

For inbound two-step transfer:

```
your_localhost% scp foo nas_username@sfel.nas.nasa.gov:foo2  
sfel% scp foo2 lou1:
```

For inbound one-step transfer if SSH-passthrough has been set up correctly:

```
your_localhost% scp foo nas_username@lou1.nas.nasa.gov:foo2
```

To transfer files through DMZFS[1,2], initiate the scp command from either a NAS HECC host or your localhost, not DMZFS[1,2]. For example,

```
your_localhost% scp foo nas_username@dmzfs1.nas.nasa.gov:foo2  
pfel% scp dmzfs1:foo2 .
```

Omit *local\_username@* and *nas\_username@* in the examples above if your local username and NAS username are identical.

◆ *Performance:*

- ◇ Within NAS HECC Enclave, depending on source and destination hosts and other factors, the performance range will be 40 - 100 Mbytes/sec.
- ◇ Over WAN (such as between NAS and a remote site), the best you get with scp from OpenSSH versions older than 4.7 (with the internal channel buffer set to 128 KB) is ~ 2 MBytes/sec. Starting with OpenSSH version 4.7, a larger channel buffer is introduced to improve file transfer performance. Users are recommended to upgrade to version 4.7 or later.

In case where OpenSSH 4.7 or a later version does not yield satisfactory performance, consider applying the [HPN-SSH patch](#) to your OpenSSH.

If the data you are transferring will compress well, consider enabling compression by adding -C to your scp command-line.

**HPN-SSH enabled scp**

HPN-SSH is a patch for OpenSSH designed to eliminate a network throughput bottleneck that typically occurs in an SSH session over long distance and high bandwidth network (i.e., when the bandwidth-delay product is high). This is

accomplished by allowing internal flow control buffers to be defined and grow at runtime, rather than statically defining them as OpenSSH does. The resulting performance increase can range from 10x to more than 50x, depending on the cipher used and host tuning.

HPN-enabled SSH is fully interoperable with other SSH servers and clients. HPN clients will be able to download faster from non-HPN servers, and HPN servers will be able to receive uploads faster from non-HPN clients. However, the host receiving the data must have a properly tuned TCP/IP stack.

Ask your local network staff for help to see if an HPN-SSH patch is needed for certain network connection.

◆ *Where is it installed at NAS?*

- ◇ On cfe2, the client version of OpenSSH 4.7p1 with HPN12v20 patch is available.
- ◇ On Lou[1-2], the client version of OpenSSH 5.0p1 with HPN13v1 patch is available.
- ◇ On SUP, both the client and server of OpenSSH 5.1p1 have been patched with HPN13v5.

On cfe2 and Lou[1-2], the HPN-patched SSH programs are purposely named as *hpn-ssh*, *hpn-scp*, and *hpn-sftp* to distinguish them from the default non-HPN versions (*ssh*, *scp* and *sftp*) of OpenSSH.

◆ *Do you need it installed on your localhost?*

To get good performance, an HPN-SSH server must be installed on your local system if data is to be received on your local system.

*Typical installation procedure:*

1. Download OpenSSH source (*openssh-x.xpx.tar.gz*) from <http://www.openssh.com>
2. Download corresponding HPN SSH patch (*openssh-x.xpx-hpnxxvx.diff.gz*) from <http://www.psc.edu/networking/projects/hpn-ssh>
3. Uncompress and untar above distributions
4. move the file *openssh-x.xpx-hpnxxvx.diff* to the directory *openssh-x.xpx*
5. cd *openssh-x.xpx* (for example, *openssh-5.0p1*)
6. `patch < openssh-5.0p1-hpn.diff`
7. `configure [OPTIONS]`
8. `make [OPTIONS]`
9. Validate:  
`%ssh -v`

◆ *Examples:*

```
lou[1-2]% hpn-scp -c arcfour -o TCPRcvBufPoll=yes source destination  
your_localhost% scp -c arcfour -o TCPRcvBufPoll=yes source destination
```

Note:

- ◇ arcfour (RC4) is a more CPU-efficient 128-bit cipher. One can also choose NONE for cipher so that there is no encryption for data.
- ◇ Enabling *TCPRcvBufPoll* allows for the TCP receive buffer to be polled through the duration of the connection.

◆ *Performance:*

With an HPN-SSH enabled scp, one can expect good performance for transferring large files to remote sites over long distance with high latency connections. Improvement over non-patched scp older than 4.7 (2 Mbytes/sec) may be 10x to 50x.

• **bbFTP**

bbFTP is a high performance remote file transfer protocol which supports parallel TCP streams for data transfers. Basically, it splits a single file in several pieces and sends them through parallel streams. The whole file is then rebuilt on the remote site. bbFTP also allows dynamically adjustable TCP/IP window sizes instead of a statically defined window size used by normal scp. In addition, it provides a secure control channel over SSH and allows data to be transferred in cleartext to reduce overhead in unnecessary encryption. These characteristics allow bbftp to achieve bandwidths that are greater than normal scp.

bbFTP is recommended in place of *scp* for the data transfer of large files over long distances.

◆ *Where is it installed at NAS?*

Both the bbFTP server (*bbftpd*) and client (*bbftp*) are installed on all Columbia hosts, Lou[1-2], Pleiades front-end and bridge nodes (pfe[1-12], bridge[1-2]) and SUP.

For DMZFS[1,2], only the bbFTP server (*bbftpd*) is installed. Issue the *bbftp* command from Columbia, Pleiades, Lou or your localhost (if bbFTP client has been installed) to push files into DMZFS[1,2] or pull files out of DMZFS[1,2].

◆ *Do you need it installed on your localhost?*

If you want to initiate *bbftp* from your localhost, you have to download and install the client version of bbFTP on your localhost. If you want to initiate *bbftp* from a NAS HECC system and transfer files from/to your localhost, download and install the server version of bbFTP on your localhost.

◆ *When to use it?*

Consider using bbFTP when transferring large files ( > 1 GB) within NAS or offsite. Be sure to use multiple streams to get better transfer rate.

◆ *Example:*

bbftp is like a non-interactive ftp and the syntax can be complicated.

```
your_localhost% bbftp -u nas_username -e 'setnbstream 8; get filename'
-E 'bbftpd -s -m 8' lou1.nas.nasa.gov
```

For formatting issue, the above command was broken into two lines. In reality, it should be just one line.

◆ *Performance:*

- ◇ bbFTP typically transfers data 10 - 20 times faster than normal scp.
- ◇ Within NAS HECC Enclave, performance should be 100+ MB/sec.
- ◇ Over WAN, the performance can be upto 50 MBytes/sec. File transfers between NAS and certain NASA sites may reach 100 Mbytes/sec.

If you are not getting good performance, check with your network administrator to see if performance tuning is needed on your system.

The article bbFTP provides more instructions on installing and using bbFTP.

• **bbSCP**

bbSCP is written in Perl by Greg Matthews at NAS. It is a bbftp wrapper which provides an scp-like command-line interface. It assembles the proper command-line for bbftp and then executes bbftp to perform the transfers. bbSCP is designed and tested for bbftp version 3.2.0.

bbSCP only encrypts usernames and passwords, it does NOT encrypt the data being transferred.

◆ *Where is it installed at NAS?*

bbSCP is installed on all Columbia hosts, Lou[1-2], Pleiades front-end and bridge nodes (pfe[1-12], bridge[1-2]) under /usr/local/bin.

◆ *Do you need it installed on your localhost?*

If you want to initiate bbSCP from your localhost, you need to:

- ◇ download and install bbftp-client-3.2.0 on your localhost

◇ download bbSCP version 1.0.6 (also attached at the end of this article)  
and install it on your localhost

◆ *When to use it?*

Use the bbSCP script when you want the bbftp functionality and performance but with scp-like syntax. It can be used for transferring files within NAS HECC Enclave or between NAS and a remote site.

◆ *Example:*

```
your_localhost% bbscp foo nas_username@lou1.nas.nasa.gov:
```

◆ *Performance:*

Performance of bbSCP is the same as bbFTP.

The article bbSCP provides more information (man page, performance turing, test and verification).

---

Article ID: 142

Last updated: 05 Aug, 2011

Data Storage & Transfer -> File Transfers -> Remote File Transfer Commands

<http://www.nas.nasa.gov/hecc/support/kb/entry/142/?ajax=1>